

Noise Folding in Compressed Sensing

Ery Arias-Castro* and Yonina C. Eldar†

April 20, 2011

Abstract

The literature on compressed sensing has focused almost entirely on settings where the signal is noiseless and the measurements are contaminated by noise. In practice, however, the signal itself is often subject to random noise prior to measurement. We briefly study this setting and show that, for the vast majority of measurement schemes employed in compressed sensing, the two models are equivalent with the important difference that the signal-to-noise ratio is divided by a factor proportional to p/n , where p is the dimension of the signal and n is the number of observations. Since p/n is often large, this leads to noise folding which can have a severe impact on the SNR.

Keywords: Compressed sensing, matching pursuit, sparse signals, analog noise vs. digital noise, noise folding.

1 Introduction

The field of compressed sensing (CS), focused on recovery of sparse vectors from few measurements, has been attracting vast interest in recent years due to its potential use in numerous signal processing applications [6, 9, 10]. The standard CS setup assumes that we are given measurements

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w}, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^n$ is the measurement vector, $\mathbf{A} \in \mathbb{R}^{n \times p}$ is the measurement matrix with $n \ll p$, and $\mathbf{w} \in \mathbb{R}^n$ is additive noise. The signal $\mathbf{x} \in \mathbb{R}^p$ is assumed to be s -sparse, so that no more than s elements of \mathbf{x} are nonzero. It is also common to assume that \mathbf{x} is deterministic, an assumption we make throughout. To recover \mathbf{x} from \mathbf{y} a variety of algorithms have been developed. These include greedy algorithms, such as thresholding and orthogonal matching pursuit (OMP) [13], and relaxation methods, such as basis pursuit [8] (also known as the Lasso) and the Dantzig selector [5].

An important aspect of CS analysis is to develop bounds on the recovery performance of these methods in the presence of noise. Two standard approaches to modeling the noise \mathbf{w} is either to assume that \mathbf{w} is deterministic and bounded [6], or that \mathbf{w} is a white noise vector, typically Gaussian [1, 2, 4, 5]. The former setting leads to a worst-case analysis in which an estimator must perform adequately even when the noise maximally damages the measurements. By contrast, if one assumes that the noise is random, then better performance bounds can be obtained. We therefore focus here on the random noise scenario.

*E. Arias-Castro is with the Department of Mathematics, University of California, San Diego, CA 92093 USA (eariasca@ucsd.edu). His work is partially supported by the Office of Naval Research under Grant N00014-09-1-0258.

†Y. C. Eldar is with the Department of Electrical Engineering, Technion Israel Institute of Technology, Haifa 32000, Israel (yonina@ee.technion.ac.il). Her work is supported in part by the Israel Science Foundation under Grant 1081/07, and by the Ollendorf foundation.

Two standard measures used to analyze the behavior of CS recovery techniques are the coherence and the restricted isometry property (RIP) [7]. If the coherence and RIP of the measurement matrix are small enough, then standard recovery methods such as basis pursuit, OMP and thresholding can recover \mathbf{x} from \mathbf{y} with squared-error that is proportional to the sparsity level s and the noise variance σ^2 , times a factor that is logarithmic in the signal length p [1, 2, 4, 5].

In many practical settings, noise is introduced to the signal \mathbf{x} prior to measurement. As an example, one of the applications of CS is to the design of sub-Nyquist A/D converters. In this context \mathbf{x} represents the analog signal at the entrance to the A/D converter, which is typically contaminated by noise [12]. Though important in practice, the prolific literature on CS has not treated signal noise in detail. Recently, several papers raised this important issue [3, 11, 14]. These works all point out the fact that noise present in \mathbf{x} can have a severe impact on the recovery performance. Here we analyze this setting in more detail in the context of the finite-CS system (1), in contrast to the analog setting treated in [3], and provide theoretical justification to these previous observations. In particular, we show that under appropriate conditions on the measurement matrix \mathbf{A} , the effect of pre-measurement noise is to degrade the signal-to-noise ratio (SNR) by a factor of n/p . In systems in which $p \ll n$ this degradation may be severe.

2 Noise Folding

2.1 Problem Formulation

The basic CS model (1) is adequate when the noise is introduced at the measurement stage, so that \mathbf{w} represents the measurement error or noise. However, in many practical scenarios, noise is added to the signal \mathbf{x} to be measured, which is not accounted for in (1). A more appropriate description in this case is

$$\mathbf{y} = \mathbf{A}(\mathbf{x} + \mathbf{z}) + \mathbf{w}, \quad (2)$$

where \mathbf{z} represents the signal noise, i.e., additive noise that is part of the signal being measured. Our goal in this letter is to analyze the effect of pre-measurement noise \mathbf{z} on the behavior of CS recovery methods. We assume throughout that \mathbf{w} is a random noise vector with covariance $\sigma^2 \mathbf{I}$, and that similarly \mathbf{z} is a random noise vector with covariance $\sigma_0^2 \mathbf{I}$, independent of \mathbf{w} . Under these assumptions, we show that (2) is equivalent to

$$\mathbf{y} = \mathbf{B}\mathbf{x} + \mathbf{u}, \quad (3)$$

where \mathbf{B} is a matrix whose coherence and RIP constants are very close to those of \mathbf{A} , and \mathbf{u} is white noise with variance $\sigma^2 + (p/n)\sigma_0^2$.

It follows that in order to study (2) we can apply the results developed for (1), with the important difference that the noise associated with (2) is larger by a factor proportional to p/n . When $n \ll p$ this leads to a large noise increase, or *noise folding*. This effect is a result of the fact that the measurement matrix \mathbf{A} aliases, or combines, all the noise elements in \mathbf{z} , even those corresponding to zero elements in \mathbf{x} , thus leading to a large noise increase in the compressed measurements.

2.2 Equivalent Formulation

To establish our results, we note that (2) can be written as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{v}, \quad (4)$$

with \mathbf{v} defined by

$$\mathbf{v} = \mathbf{w} + \mathbf{A}\mathbf{z}. \quad (5)$$

Under the assumption of white noise, the effective noise vector \mathbf{v} has covariance \mathbf{Q} , where

$$\mathbf{Q} = \sigma^2 \mathbf{I} + \sigma_0^2 \mathbf{A} \mathbf{A}^T. \quad (6)$$

As can be seen, in general \mathbf{v} is no longer white, which complicates the recovery analysis.

A simple special case is when $\mathbf{A} \mathbf{A}^T$ is proportional to the identity, so that \mathbf{v} is still white noise. As an example, suppose that \mathbf{A} is the concatenation of $r = p/n$ orthonormal bases, i.e., $\mathbf{A} = [\mathbf{A}_1 \cdots \mathbf{A}_r]$, where each \mathbf{A}_k is an $n \times n$ orthogonal matrix—for example, we may want to analyze a signal with a few bases (e.g, wavelets and sinusoids) as in [8]. In this case,

$$\mathbf{A} \mathbf{A}^T = \mathbf{A}_1 \mathbf{A}_1^T + \cdots + \mathbf{A}_r \mathbf{A}_r^T = r \mathbf{I} = \frac{p}{n} \mathbf{I}$$

so that the noise covariance (6) becomes $\mathbf{Q} = \gamma \mathbf{I}$ with

$$\gamma = \sigma^2 + \frac{p}{n} \sigma_0^2. \quad (7)$$

In this special case the models of (4) (or (2)) and (1) are identical, with the only difference being that the noise variance of \mathbf{v} has increased by γ/σ^2 with respect to that of \mathbf{w} . Assuming that $\sigma_0 \approx \sigma^2$ the increase in noise is proportional to p/n , a simple case of *noise folding*.

In the next section we show that this result holds more generally. Namely, the models of (4) and (1) are roughly equivalent with a noise increase of γ/σ^2 even when $\mathbf{A} \mathbf{A}^T$ is not proportional to the identity.

3 RIP and Coherence with Whitening

Consider now a more general CS setting, where \mathbf{A} is an arbitrary matrix with low coherence or low RIP. To proceed, we first whiten the noise \mathbf{v} in (4) by multiplying the linear system by $\mathbf{Q}_1^{-1/2}$, where $\mathbf{Q}_1 := \mathbf{Q}/\gamma$, obtaining the equivalent system

$$\mathbf{y} = \mathbf{B} \mathbf{x} + \mathbf{u}, \quad \mathbf{B} := \mathbf{Q}_1^{-1/2} \mathbf{A}, \mathbf{u} := \mathbf{Q}_1^{-1/2} \mathbf{v}. \quad (8)$$

Now, the noise vector \mathbf{u} is white with covariance matrix $\gamma \mathbf{I}$, just as in the case of $\mathbf{A} \mathbf{A}^T$ proportional to the identity. The main difference, however, is that the whitening changed the measurement matrix from \mathbf{A} to \mathbf{B} . We quantify the magnitude of these changes below via the RIP constants and the coherence. As we show, for standard matrices used in CS, the change is generally not very significant.

Our results hinge on approximating $\mathbf{A} \mathbf{A}^T$ by $(p/n) \mathbf{I}$ even when \mathbf{A} is arbitrary. Let

$$\eta := \|\mathbf{I} - (n/p) \mathbf{A} \mathbf{A}^T\|_2, \quad (9)$$

measure the quality of this approximation, where $\|\cdot\|_2$ denotes the standard operator norm in \mathbb{R}^n . In our derivations below we will assume that η is small. Under this assumption we will show that the coherence and RIP constants of \mathbf{B} and \mathbf{A} are very similar.

To justify the assumption that η is often small, note that when the entries of \mathbf{A} are i.i.d. zero-mean, variance $1/n$ random variables with a sub-gaussian distribution (including the normal, uniform, Bernoulli distributions), or when the column vectors are i.i.d. uniform on the sphere, then $\eta \leq C \sqrt{n/p}$ with probability at least $1 - \exp(-cn)$, for constants $C, c > 0$ depending only on the distribution of \mathbf{A} [15, Th. 39]. For example, when the entries of \mathbf{A} are i.i.d. $N(0, 1/n)$ and n is large enough,

$$\eta \leq 2\sqrt{n/p} + n/p + 4t/\sqrt{p}, \quad (10)$$

with probability at least $1 - 2\exp(-t^2/2)$ if $0 < t \leq \sqrt{n}$ and $p \geq n$ [15, Cor. 35]. A similar result holds for other distributions, including heavy-tailed distributions, the basic requirement being that the column vectors of \mathbf{A} are independent with covariance \mathbf{I}/n [15]. These assumptions are standard in the CS literature. Thus, in these prevalent settings, η will be small with high probability.

3.1 Restricted Isometry Analysis

We begin by showing that the restricted isometry constants of \mathbf{B} and \mathbf{A} are similar assuming a small value of η .

For an index set $\Lambda \subset \{1, \dots, p\}$ of size s , let \mathbf{A}_Λ denote the submatrix of \mathbf{A} made of the column vectors indexed by Λ . We say that \mathbf{A} has the RIP with constants $0 < \alpha_s \leq \beta_s$ if

$$\alpha_s \|\mathbf{h}\|^2 \leq \|\mathbf{A}_\Lambda \mathbf{h}\|^2 \leq \beta_s \|\mathbf{h}\|^2, \quad \forall \mathbf{h} \in \mathbb{R}^s, \quad (11)$$

for any index set $\Lambda \subset \{1, \dots, p\}$ of size s . The following proposition relates the RIP constants of \mathbf{B} and \mathbf{A} :

Proposition 1. *Assume that $\eta < 1/2$ in (9) and that \mathbf{A} satisfies the RIP of order s with constants $0 < \alpha_s \leq \beta_s$. Then \mathbf{B} satisfies the RIP of order s with constants $\alpha_s(1 - \eta_1)$ and $\beta_s(1 + \eta_1)$, where $\eta_1 := \eta/(1 - \eta)$.*

Though the bound is valid for $\eta < 1$, the smaller RIP constant for \mathbf{B} is only positive when $\eta < 1/2$, leading to our restriction.

Proof. The proof is based on the fact that \mathbf{Q}_1 is close to \mathbf{I} . Indeed, by definition of η in (9),

$$\begin{aligned} \|\mathbf{Q}_1 - \mathbf{I}\|_2 &= \frac{\sigma_0^2}{\gamma} \|\mathbf{A}\mathbf{A}^T - (p/n)\mathbf{I}\|_2 \\ &= \frac{\sigma_0^2(p/n)}{\sigma^2 + \sigma_0^2(p/n)} \eta \leq \eta. \end{aligned}$$

Next, we express $\mathbf{Q}_1^{-1} - \mathbf{I}$ as a power series

$$\mathbf{Q}_1^{-1} - \mathbf{I} = (\mathbf{I} - (\mathbf{I} - \mathbf{Q}_1))^{-1} - \mathbf{I} = \sum_{k \geq 1} (\mathbf{I} - \mathbf{Q}_1)^k,$$

which converges since $\|\mathbf{Q}_1^{-1} - \mathbf{I}\|_2 \leq \eta < 1$ and $\|\cdot\|_2$ is an operator norm. Taking norms on both sides of the inequality and using both the triangle inequality and again the fact that $\|\cdot\|_2$ is an operator norm, we obtain

$$\begin{aligned} \|\mathbf{Q}_1^{-1} - \mathbf{I}\|_2 &\leq \sum_{k \geq 1} \|\mathbf{Q}_1 - \mathbf{I}\|_2^k \\ &\leq \sum_{k \geq 1} \eta^k = \frac{\eta}{1 - \eta} = \eta_1. \end{aligned} \quad (12)$$

Let Λ be an index set of size s and take any $\mathbf{h} \in \mathbb{R}^s$. Then,

$$\|\mathbf{B}_\Lambda \mathbf{h}\|^2 - \|\mathbf{A}_\Lambda \mathbf{h}\|^2 = \mathbf{h}^T \mathbf{A}_\Lambda^T (\mathbf{Q}_1^{-1} - \mathbf{I}) \mathbf{A}_\Lambda \mathbf{h}.$$

Since

$$\begin{aligned} |\mathbf{h}^T \mathbf{A}_\Lambda^T (\mathbf{Q}_1^{-1} - \mathbf{I}) \mathbf{A}_\Lambda \mathbf{h}| &\leq \|\mathbf{Q}_1^{-1} - \mathbf{I}\|_2 \|\mathbf{A}_\Lambda \mathbf{h}\|^2 \\ &\leq \eta_1 \|\mathbf{A}_\Lambda \mathbf{h}\|^2, \end{aligned}$$

we have that

$$(1 - \eta_1) \|\mathbf{A}_\Lambda \mathbf{h}\|^2 \leq \|\mathbf{B}_\Lambda \mathbf{h}\|^2 \leq (1 + \eta_1) \|\mathbf{A}_\Lambda \mathbf{h}\|^2.$$

Together with (11), we obtain

$$\alpha_s(1 - \eta_1) \|\mathbf{h}\|^2 \leq \|\mathbf{B}_\Lambda \mathbf{h}\|^2 \leq \beta_s(1 + \eta_1) \|\mathbf{h}\|^2,$$

which concludes the proof. \square

3.2 Coherence Analysis

We now turn to analyze the coherence. Denoting by \mathbf{A}_i the i th column vector of \mathbf{A} , the coherence of \mathbf{A} is defined as

$$\mu(\mathbf{A}) = \max_{i \neq j} \frac{|\mathbf{A}_i^T \mathbf{A}_j|}{\|\mathbf{A}_i\| \|\mathbf{A}_j\|}.$$

Proposition 2. Assume that $\eta < 3/4$ in (9). Then

$$\mu(\mathbf{B}) \leq (1 + \eta_2) \mu(\mathbf{A}),$$

where

$$\eta_2 := (2\sqrt{1 - \eta} - 1)^{-2} - 1.$$

Note that $\eta_2 = 2\eta + O(\eta^2)$, with $\eta_2 < 5\eta$ when $\eta < 1/2$.

Proof. To prove the proposition we develop an upper bound on the numerator $|\mathbf{B}_i^T \mathbf{B}_j|$ of $\mu(\mathbf{B})$, and a lower bound on the denominator elements $\|\mathbf{B}_i\|$. For $i \neq j$, we have

$$\begin{aligned} |\mathbf{B}_i^T \mathbf{B}_j| &= |\mathbf{A}_i^T \mathbf{Q}_1^{-1} \mathbf{A}_j| \\ &\leq |\mathbf{A}_i^T \mathbf{A}_j| + |\mathbf{A}_i^T (\mathbf{Q}_1^{-1} - \mathbf{I}) \mathbf{A}_j| \\ &\leq (1 + \eta_1) |\mathbf{A}_i^T \mathbf{A}_j|, \end{aligned}$$

by (12).

We now lower bound $\|\mathbf{B}_i\|$ in terms of $\|\mathbf{A}_i\|$ and η . In parallel with the proof of Proposition 1, we express $\mathbf{Q}_1^{-1/2} - \mathbf{I}$ as a power series

$$\mathbf{Q}_1^{-1/2} - \mathbf{I} = \sum_{k \geq 1} c_k (\mathbf{I} - \mathbf{Q}_1)^k,$$

where c_k are the coefficients in the Taylor expansion of $(1 - x)^{-1/2}$. Taking norms on both sides of the inequality, we obtain

$$\begin{aligned} \|\mathbf{Q}_1^{-1/2} - \mathbf{I}\|_2 &\leq \sum_{k \geq 1} c_k \|\mathbf{Q}_1 - \mathbf{I}\|_2^k \\ &\leq \sum_{k \geq 1} c_k \eta^k = (1 - \eta)^{-1/2} - 1. \end{aligned}$$

Therefore,

$$\begin{aligned} \|\mathbf{B}_i\| &= \|\mathbf{Q}_1^{-1/2} \mathbf{A}_i\| \\ &\geq \|\mathbf{A}_i\| - \|(\mathbf{Q}_1^{-1/2} - \mathbf{I}) \mathbf{A}_i\| \\ &\geq (1 - \eta_3) \|\mathbf{A}_i\|, \end{aligned}$$

where $\eta_3 = (1 - \eta)^{-1/2} - 1$. All together, we have

$$\frac{|\mathbf{B}_i^T \mathbf{B}_j|}{\|\mathbf{B}_i\| \|\mathbf{B}_j\|} \leq \frac{(1 + \eta_1) |\mathbf{A}_i^T \mathbf{A}_j|}{(1 - \eta_3)^2 \|\mathbf{A}_i\| \|\mathbf{A}_j\|},$$

with $(1 + \eta_1)/(1 - \eta_3)^2 = 1 + \eta_2$ by definition of η_2 . \square

4 Conclusion

Though the CS literature is almost silent on the effect of pre-measurement noise on recovery performance, in this letter we made the point that it may have a substantial impact on SNR. Indeed, we showed that, for the most common measuring schemes used in CS, the model with pre-measurement noise is, after whitening, equivalent to a standard model with only measurement noise, modulo a change in measurement matrix and an increase in the noise variance by a factor of p/n . We provided rigorous bounds on the RIP constants and the coherence of the new measurement matrix which show that, as $n, p \rightarrow \infty$ with $p/n \rightarrow 0$, the constants are essentially unchanged. As the performance of standard recovery methods are formulated in terms of the RIP constants and the coherence, this shows that, in this regime, these methods operate as usual, except with noise folding lead to a noise variance multiplied by p/n .

References

- [1] Z. Ben-Haim and Y. C. Eldar. The cramér-rao bound for estimating a sparse parameter vector. *IEEE Trans. Signal Processing*, 58:3384–3389, June 2010.
- [2] Z. Ben-Haim, Y. C. Eldar, and M. Elad. Coherence-based performance guarantees for estimating a sparse vector under random noise. *Trans. Sig. Proc.*, 58:5030–5043, October 2010.
- [3] Z. Ben-Haim, T. Michaeli, and Y. C. Eldar. Performance bounds and design criteria for estimating finite rate of innovation signals. 2010. submitted to *IEEE. Trans. on Info. Theory*.
- [4] P. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- [5] E. Candès and T. Tao. The Dantzig Selector: Statistical estimation when p is much larger than n . 35(6):2313–2351, 2007.
- [6] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59:1207–1223, 2006.
- [7] E. J. Candès and M. B. Wakin. An introduction to compressive sampling [a sensing/sampling paradigm that goes against the common knowledge in data acquisition]. *IEEE Signal Processing Magazine*, 25(2):21–30, March 2008.
- [8] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.*, 20(1):33–61, 1998.
- [9] D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.
- [10] M. F. Duarte and Y. C. Eldar. Structured compressed sensing: Theory and applications. 2010. Submitted.

- [11] B. Miller, J. Goodman, K. Forsythe, J. Sun, and V. Goyal. A multi-sensor compressed sensing receiver: Performance bounds and simulated results. In *Signals, Systems and Computers, 2009 Conference Record of the Forty-Third Asilomar Conference on*, pages 1571–1575, 2009.
- [12] M. Mishali and Y. C. Eldar. Xampling: Compressed sensing for analog signals. In Y. C. Eldar and G. Kutyniok, editors, *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2011.
- [13] Y. Pati, R. Rezaifar, and P. Krishnaprasad. Orthogonal Matching Pursuit: Recursive function approximation with applications to wavelet decomposition. Pacific Grove, CA, Nov. 1993.
- [14] J. Treichler, M. Davenport, and R. Baraniuk. Application of compressive sensing to the design of wideband signal acquisition receivers. In *Proc. 6th U.S. / Australia Joint Workshop on Defense Applications of Signal Processing (DASP)*, Lihue, Hawaii, 2009.
- [15] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. Available from <http://arxiv.org/abs/1011.3027>, 2010.